

The Ditmarsch Tale of Wonders — The Dynamics of Lying

Hans van Ditmarsch, University of Seville, hvd@us.es

Abstract

We propose a dynamic logic of lying, wherein a ‘lie that φ ’ (where φ is a formula in the logic) is an action in the sense of dynamic modal logic, that is interpreted as a state transformer relative to the formula φ . The states that are being transformed are pointed Kripke models encoding the uncertainty of agents about their beliefs. Lies can be about factual propositions but also about modal formulas, such as the beliefs of other agents or the belief consequences of the lies of other agents. We distinguish (i) an outside observer who is lying to an agent that is modelled in the system, from (ii) one agent who is lying to another agent, and where both are modelled in the system. For either case, we further distinguish (iii) the agent who believes everything that it is told (even at the price of inconsistency), from (iv) the agent who only believes what it is told if that is consistent with its current beliefs, and from (v) the agent who believes everything that it is told by consistently revising its current beliefs. The logics have complete axiomatizations, which can most elegantly be shown by way of their embedding in what is known as action model logic or the extension of that logic to belief revision.

1 Introduction

My favourite of Grimm’s fairytales is ‘Hans im Glück’ (Hans in Luck). A close second comes ‘The Ditmarsch Tale of Wonders’. In German this is called a ‘Lügenmärchen’, a ‘Liar’s Tale’. It is as follows.

I will tell you something. I saw two roasted fowls flying; they flew quickly and had their breasts turned to Heaven and their backs to Hell; and an anvil and a mill-stone swam across the Rhine prettily, slowly, and gently; and a frog sat on the ice at Whitsuntide and ate a ploughshare.

Four fellows who wanted to catch a hare, went on crutches and stilts; one of them was deaf, the second blind, the third dumb, and the fourth could not stir a step. Do you want to know how it was done? First, the blind man saw the hare running across the field, the dumb one called to the deaf one, and the lame one seized it by the neck.

There were certain men who wished to sail on dry land, and they set their sails in the wind, and sailed away over great fields. Then they sailed over a high mountain, and there they were miserably drowned.

A crab was chasing a hare which was running away at full speed; and high up on the roof lay a cow which had climbed up there. In that country the flies are as big as the goats are here.

Open the window that the lies may fly out. [19]

A passage like “A crab was chasing a hare which was running away at full speed; and high up on the roof lay a cow which had climbed up there.” contains very obvious lies. Nobody considers it possible that this is true. Crabs are reputedly slow, hares are reputedly fast.

In ‘The Ditmarsch Tale of Wonders’, none of the lies are believed.

In the movie ‘The Invention of Lying’¹ the main character Mark goes to a bank counter and finds out he only has \$300 in his account. But he needs \$800. Lying has not yet been invented in the 20th-century fairytale country of this movie — that however seems to be in a universe very parallel to either the British Midlands or Brooklyn New York. Then and there, on the spot, Mark invents lying. We see some close-ups of Mark’s braincells doing heavy duty—such a thing has not happened before. And then, Mark tells the bank employee assisting him that there must be a mistake: he has \$800 in his account. He is lying. She responds, oh well, then there must be a mistake with your account data, because on my screen it says you only have \$300. I’ll inform system maintenance of the error. My apologies for the inconvenience. And she gives him \$800! In the remainder of the movie, Mark gets very rich.

Mark’s lies are not as unbelievable as those in Grimm’s fairytale. It is possible that he has \$800. It is just not true. Still, there is something unrealistic about the lies in this movie: new information is believed instantly. New information is even believed if it is inconsistent with prior information. After Mark’s invention of lying while obtaining \$800, he’s trying out his invention on many other people. It works all the time! There are shots wherein he first announces a fact, then its negation, then the fact again, while all the time his extremely credulous listeners keep believing every last announcement. New information is also believed if it contradicts direct observation. In a café, in company of several of his friends, he claims to be a one-armed bandit. And they commiserate with him, oh, I never knew you only had one arm, how terrible for you. All the time, Mark is sitting there drinking beer and gesturing with both hands while telling his story.

In the movie ‘The Invention of Lying’, all lies are believed.

¹http://en.wikipedia.org/wiki/The_Invention_of_Lying

In the real world, if you lie, sometimes other people believe you and sometimes they don't. When can you get away with a lie? Consider the consecutive numbers riddle [25].

Anne and Bill are each going to be told a natural number. Their numbers will be one apart. The numbers are now being whispered in their respective ears. They are aware of this scenario. Suppose Anne is told 2 and Bill is told 3.

The following truthful conversation between Anne and Bill now takes place:

- *Anne: "I do not know your number."*
- *Bill: "I do not know your number."*
- *Anne: "I know your number."*
- *Bill: "I know your number."*

Explain why is this possible.

Initially, Anne is uncertain between Bill having 1 or 3, and Bill is uncertain between Anne having 2 or 4. So both Anne and Bill do not initially know their number.

Suppose that Anne first says to Bill: "I know your number." Anne is lying. Bill does not consider it possible that Anne knows his number. However, Anne did not know that Bill would not believe her. She considered it possible that Bill had 1, in which case Bill would have considered it possible that Anne was telling the truth, and would then have drawn the incorrect conclusion that Anne had 0.

Alternatively, suppose that the first announcement, Anne's, is truthful, but that Bill is lying in the second announcement and says to Anne: "I know your number." If Anne believes that, she will then say "I know your number," as she believes Bill to have 1. Her announcement is an honest mistake, because this belief is incorrect. However, as a result of Anne's announcement Bill will learn Anne's number, so that his announcement "I know your number," that was a lie at the time, now has become true.

That is, if you are still following us.

It seems not so clear how all this should be formalized in a logic interpreted on epistemic modal structures, and this is the topic of our paper.

In real life, some lies are believed and some are not.

1.1 The modal dynamics of lying

What is a lie? Let p be a proposition. You lie to me that p , if you believe that p is false while you say that p , and with the intention that I believe p . The thing you say we call the announcement. If you succeed in your intention, I believe p , and I also believe that your announcement of p was truthful, i.e., that you believed that p when you said that p . In this investigation we abstract from the intentional aspect of lying. Such an abstraction seems reasonable. It is similar to that in AGM belief revision [2], wherein one models how to incorporate new information in an agent's belief set, but abstracts from the process that made the new information acceptable to the agent. Our proposal is firmly grounded in modal logic. We employ dynamic epistemic logic [44].

Preconditions What are the modal preconditions and postconditions of a lie? Let us for now assume that p itself is not a modal proposition but a Boolean proposition. We further assume two agents, a and b . Agent a will be assumed female and agent b will be assumed male. Typically, in our exposition a will be the speaker or sender and b will be the receiver or addressee. However, a and b are not agent roles but agent names (and agent variables). We also model dialogue wherein agents speak in turn; so these roles may swap. Formula B_ap stand for ‘agent a believes that p ’. We use belief modalities B and not knowledge modalities K , because lying results in false beliefs, whereas knowledge modalities are used for correct beliefs.

The precondition of ‘ a is lying that p to b ’ is $B_a\neg p$ (\neg is negation). Stronger preconditions are conceivable, e.g., that the addressee considers it possible that the lie is true, $\neg B_b\neg p$, or that the speaker believes that, $B_a\neg B_b\neg p$. These conditions may not always hold while we still call the announcement a lie, because the speaker may not know whether the additional conditions are satisfied. We therefore will only require precondition $B_a\neg p$.

We should contrast the announcement that p by a lying agent with other forms of announcement. Just as a lying agent believes that p is false when it announces p , a truthful agent believes that p is true when it announces p . The precondition for a lying announcement by a is $B_a\neg p$, and so the precondition for a truthful announcement by a is B_ap . Here, we should put up some strong terminological barriers in order to avoid pitfalls and digressions into philosophy and epistemology. Truthful is synonymous with honest. Now dictionaries, that report actual usage, do not make a difference between an agent a telling the truth and an agent a believing that she is telling the truth. A modal logician has to make a choice. We mean the latter, exclusively. A truthful announcement may therefore not be a true announcement. If p is false but agent a mistakenly believes that p , then when she says p , that is a truthful but false announcement of p . Besides the truthful and the lying announcement there is yet another form of announcement, because in modal logic there are always three instead of two possibilities: either you believe p , or you believe $\neg p$, or you are uncertain whether p . The last corresponds to the precondition $\neg(B_ap \vee B_a\neg p)$ (\vee is disjunction). An announcement wherein agent a announces p while she is uncertain about p we propose to call a bluffing announcement. The dictionary meaning for the verb bluff is ‘to cause to believe what is untrue’ or ‘to deceive or to feign’ in a more general sense. Its meaning is even more intentional than that of lying. Feigning belief in p means suggesting belief in p , by saying it, or otherwise behaving in accordance to it, although you do not have this belief. This corresponds to $\neg B_ap$ as precondition. This would make lying a form of bluffing, as $B_a\neg p$ implies $\neg B_ap$. It is common and according to Gricean conversational norms to consider that saying something that you believe to be false is worse than (or, at least, different from) saying something that you do not believe to be true. This brings us to $\neg B_ap \wedge \neg B_a\neg p$ (\wedge is conjunction), equivalent to $\neg(B_ap \vee B_a\neg p)$.

To the three mutually exclusive (and complete) preconditions B_ap , $B_a\neg p$, and $\neg(B_ap \vee B_a\neg p)$ we associate the truthful, lying, and bluffing announcement that p , and we call the announcing agent a truth teller, liar and bluffer, respectively.² The three forms of

²Here again, it is stretching usage that a truth teller may not be telling the truth but only what she

announcement are intricately intertwined. This is obvious: to a credulous addressee a lying announcement appears to be a truthful announcement, whereas a skeptical addressee, who already believed the opposite of the announcement, has to make up his mind whether the speaker is merely mistaken, or is bluffing, or is even lying.

Postconditions We now consider the postconditions of ‘ a is lying that p to b ’. If a ’s intention to deceive b is successful, b believes p after the lie. Therefore, $B_b p$ should be a postcondition of a successful execution of the action of lying. Also, the precondition should be preserved: $B_a \neg p$ should still true after the lie. In the first place, we propose logics to achieve this. However, this comes at a price. In case the agent b already believed the opposite, $B_b \neg p$, then b ’s beliefs are inconsistent afterwards. (This merely means that b ’s accessibility relation is empty, not that the logic is inconsistent.) There are two different solutions for this: either b does not change his beliefs, so that $B_b \neg p$ still holds after the lie, or the belief $B_b \neg p$ is given up in order to consistently incorporate $B_b p$. The three alternative postconditions after the lie that p are therefore: (i) always make $B_b p$ true after the lie (even at the price of inconsistency), (ii) only make $B_b p$ true if agent b considered p possible before the lie ($\neg B_b \neg p$), and (iii) always make $B_b p$ true by a consistency preserving process of belief revision. These are all modelled.

Lying as a dynamic modal operator The preconditions and postconditions of lying may contain epistemic modal operators (belief modalities). The action of lying itself is modelled as a dynamic modal operator. The dynamic modal operator for ‘lying that p ’ is interpreted as an epistemic state transformer. An epistemic state is a pointed Kripke model (a model with a designated state) that encodes the beliefs of the agents. An epistemic action ‘agent a lies that p to agent b ’ should transform an epistemic state satisfying $B_a \neg p$ into an epistemic state satisfying $B_b p$ and $B_a \neg p$. The execution of such dynamic modal operators for epistemic actions depends on the initial epistemic state and that operator’s description only. In that sense, they are different from dynamic modal operators for (PDL-style) actions that are interpreted using an accessibility relation in a given structure.

In this dynamic epistemic setting we can distinguish (i) the case of an external observer (an agent who is not explicitly modelled in the structures and in the logical language), who is lying to an agent modelled in the system, from (ii) the case of one agent lying to another agent, where both are explicitly modelled. For this external agent a truthful announcement is the same as a true announcement, and a lying announcement is the same as a false announcement.³ These matters will be addressed in detail.

In dynamic epistemic logics the transmission of messages is instantaneous and infallible. This is another assumption in our modelling framework.

believes to be the truth, but that cannot be helped.

³This explains the terminological confusion in the area: the logic known as that of truthful public announcements is really the logic of *true* public announcements.

Lying about modal formulas The belief operators B_a do not merely apply to Boolean propositions p but to any proposition φ with belief modalities. This is known as higher-order belief. In the semantics, the generalization from ‘lying that p ’ to ‘lying that φ ’ for any proposition, does not present any problem. This is, because it will be defined relative to the set of states where the formula is believed by the speaker a . For a subset of the domain, it does not matter if it was determined for a Boolean formula or for a modal formula. Still, there are other problems.

Firstly, we aim for agents having what is known as ‘normal’ beliefs, that satisfy consistency and introspection: $B_a\varphi \rightarrow \neg B_a\neg\varphi$, $B_a\varphi \rightarrow B_aB_a\varphi$, and $\neg B_a\varphi \rightarrow B_a\neg B_a\varphi$. If we wish the addressee b to believe the lie that p even if he already believed $\neg p$, his beliefs would become inconsistent. The property $B_a\varphi \rightarrow \neg B_a\neg\varphi$ of consistent belief is therefore not preserved under lying updates. We will address this.

Secondly, we may well require that $B_a\neg p$ is true before the lie by a that p and that $B_a\neg p$ and B_bp are true after the lie that p , but we cannot and do not even want to require for any φ that, if $B_a\neg\varphi$ is true before the lie by a that φ , then $B_a\neg\varphi$ and $B_b\varphi$ are true afterwards. For a typical example, suppose that p is false and that a lies that $p \wedge \neg B_bp$ to b .⁴ We would say that the lie was successful if B_bp holds, not if $B_b(p \wedge \neg B_bp)$ holds, an inconsistency (for belief). Also, we do not want that $B_a\neg(p \wedge \neg B_bp)$, that was true before the lie, remains true after the lie. It is common in the area to stick to the chosen semantic operation but not to require persistence of belief in such cases.

Our proposed modelling, that also applies to modal formulas, allows us to elegantly explain why in the consecutive number riddle we can have that (see above) ‘*as a result of Anne’s announcement Bill will learn Anne’s number, so that his prior announcement “I know your number,” that was a lie at the time, now has become true.*’ The seemingly contradictory announcements in the riddle are Moorean phenomena, and the added aspect of uncertainty about lying or truth-telling makes the analysis more complex, and the results more interesting.

More modalities In the concluding Section 9 we discuss further issues in the modal logic of lying, such as group epistemic operators (common belief) in preconditions and postconditions of lying, and structures with histories of actions to keep track of the number of past lies.

1.2 A short history of lying

We conclude this introduction with a review of literature on lying.

Philosophy Lying has been a thriving topic in the philosophical community for a long, long time [36, 11, 28, 29]. Almost any analysis starts with quoting Augustine on lying:

⁴It is not raining in Sevilla. You don’t know that. I am lying to you: “You don’t know that it is raining in Sevilla!” I.e., using the conventional conversational implicature, “It is raining in Sevilla but/and you do not know that.” This is a Moorean sentence.

“that man lies, who has one thing in his mind and utters another in words”
 “the fault of him who lies, is the desire of deceiving in the uttering of his mind”
 [5]

In other words: saying that p while believing that $\neg p$, with the intention to make believe p , our starting assumption. The requirements for the belief preconditions and postconditions in such works are illuminating [29]. For example, the addressee should not merely believe the lie but believe it to be believed by the speaker. Indeed, ... and even believed to be commonly believed, would the modal logician say (see the final Section 9). Scenarios involving eavesdroppers (can you lie to an agent who is not the addressee?) are relevant for logic and multi-agent system design, and also claims that you can only lie if you really say something: an omission is not a lie [29]. Wrong, says the computer scientist: if the protocol is common knowledge, you can lie by not acting when you should have; say, by not stepping forward in the muddy children problem, although you know that you are muddy. The philosophical literature also clearly distinguishes between false propositions and propositions believed to be false but in fact true, so that when you lie about them, in fact you tell the truth. Gettier-like scenarios are presented, including delayed justification [34].⁵ Much is said on the morality of lying [11] and on its intentional aspect. As said, we abstract from the intentional aspect of lying. We also abstract from its moral aspect.

Psychology Lying excites great interest in the general public. Lots of popular science books are written on the topic, typical examples are [38, 41]. In psychology, biology, and other experimental sciences lying and deception are related. A cuckoo is ‘lying’ if it is laying eggs in another bird’s nest. Two issues are relevant for our investigation. Firstly, that it is typical to be believed, and that lying is therefore the exception. We model the ‘successful’ lie that is indeed believed, unless there is evidence to the contrary: prior belief in the opposite. Secondly, that the detection of lying is costly, and that this is a reason to be typically believed. In logic, cost is computational complexity. The issue of the complexity of lying is shortly addressed in Section 9 on further research.

Economics In economics, ‘cheap talk’ is making false promises. Your talk is cheap if you do not intend to execute an action that you publicly announced to plan. It is therefore a lie, it is deception [18, 22]. Our focus is different. We do not model lying about planned actions but lying about propositions, and in particular on their belief consequences. Economists postulate probabilities for lying strategies and truthful strategies, to be tested experimentally. We only distinguish lies that are always believed from lies that (in the face of contradictory prior belief) are never believed.

Logic Papers that model lying as an epistemic action, inducing a transformation of an epistemic model, include [6, 37, 8, 43, 24, 45]. Lying by an external observer has been discussed by Baltag and collaborators from the inception of dynamic epistemic logic onward

⁵Suppose that you believe that $\neg p$ and that you lie that p . Later find out that your belief was mistaken because p was really true. You can then with some justification say “Ah, so I was not really lying.”

[6]; the later [8] also discusses lying in logics with plausible belief, as does [43]. In [45] the conscious update in [17] is applied to model lying by an external observer. In [35] the authors give a modal logic of lying and bluffing, including intentions. Instead of bluffing they call this bullshit, after [16]. Strangely, in view of contradictory Moorean phenomena, they do not model lying as a dynamic modality. In [37, 24] the unbelievable update is considered; this is the issue consistency preservation for belief, as in our treatment of unbelievable lies (rejecting the lie that p if you already believe $\neg p$). The promising manuscript [26] allows explicit reference in the logical language to truthful, lying and bluffing agents (the authors call this ‘agent types’), thus enabling some form of self-reference.

Artificial intelligence? Various of the already cited works could have been put under this header. But then it would not have been a question mark. Applications of epistemic logic in artificial intelligence typically are about knowledge and not about belief. Successful frameworks as interpreted systems and knowledge programs [15] model multi- $\mathcal{S5}$ systems. Our analysis of the modal dynamics of lying aims to prepare the ground for AI applications involving belief instead of knowledge, and to facilitate determining the complexities of such reasoning tasks.⁶

1.3 Contributions and overview

A main and novel contribution of our paper is a precise model of the informative consequences of two agents lying to each other, and a logic for that, including a treatment of bluffing. This agent-to-agent-lying, in the logic called agent announcement logic, is presented in Section 4, with an extended example in Section 5. A special, simpler, case is that of an outside observer who is lying to an agent that is modelled in the system. This (truthful and lying) public announcement logic is treated in Section 3. That section mainly contains results from [45]. Section 2 introduces the standard truthful (without lying) public announcement logic of which our proposals can be seen as variations. Section 6 on action models is an alternative perspective on the frameworks presented in Section 3 and Section 4. It anchors them in another part of dynamic epistemic logic. Section 7 contains another novel contribution. It adapts the logics of the Sections 3 and 4 to the requirement that unbelievable lies (if you hear p but already believe $\neg p$) should not be incorporated. Subsequently, Section 8 adapts these logics to the requirement that unbelievable lies, on the contrary, should be incorporated, but consistently so. This can be anchored in yet another part of the dynamic epistemic logical literature, involving structures with plausibility relations. All these logics have complete axiomatizations (which is unremarkable). An incidental novel contribution is how to resolve ambiguity between bluffing and lying with disjunctive normal forms (Proposition 18). Section 9 finishes the paper with considerations on the limitations of our approach and further research.

⁶Argumentation theory, when seen as an area in AI, does of course model beliefs and their justifications.

2 Truthful public announcements

The well-known logic of truthful public announcements [33, 7] is an extension of multi-agent epistemic logic. Its language, structures, and semantics are as follows. Given are a finite set of agents A and a countable set of propositional variables P .

Definition 1 The *language of truthful public announcement logic* is inductively defined as

$$\mathcal{L}(!) \ni \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \psi) \mid B_a\varphi \mid [!\varphi]\psi$$

where $p \in P$, and $a \in A$.⁷ Without the announcement operators we get the language \mathcal{L} of epistemic logic. \dashv

Other propositional connectives are defined by abbreviation. For $B_a\varphi$, read ‘agent a believes formula φ ’. Agent variables are a, b, c, \dots . For $[!\varphi]\psi$, read ‘after truthful public announcement of φ , formula ψ (is true)’. The dual operator for the necessity-type announcement operator is by abbreviation defined as $\langle !\varphi \rangle\psi := \neg[!\varphi]\neg\psi$. If $B_a\neg\varphi$ we say that φ is *unbelievable* and, consequently, if $\neg B_a\neg\varphi$ we say that φ is *believable*. This is also read as ‘agent a considers it possible that φ ’.

Definition 2 An *epistemic model* $M = (S, R, V)$ consists of a *domain* S of *states* (or ‘worlds’), an *accessibility function* $R : A \rightarrow \mathcal{P}(S \times S)$, where each $R(a)$, for which we write R_a , is an accessibility relation, and a *valuation* $V : P \rightarrow \mathcal{P}(S)$, where each $V(p)$ represents the set of states where p is true. For $s \in S$, (M, s) is an *epistemic state*. \dashv

An epistemic state is also known as a pointed Kripke model. We often omit the parentheses. Four model classes will appear in this work. Without any restrictions we call the model class \mathcal{K} . The class of models where all accessibility relations are transitive and euclidean is called $\mathcal{K}45$, and if they are also serial it is called $\mathcal{KD}45$. The class of models where all accessibility relations are equivalence relations is $\mathcal{S}5$. Class $\mathcal{KD}45$ is said to have the *properties of belief*, and $\mathcal{S}5$ to have the *properties of knowledge*.

Definition 3 Assume an epistemic model $M = (S, R, V)$.

$$\begin{aligned} M, s &\models p && \text{iff } s \in V_p \\ M, s &\models \neg\varphi && \text{iff } M, s \not\models \varphi \\ M, s &\models \varphi \wedge \psi && \text{iff } M, s \models \varphi \text{ and } M, s \models \psi \\ M, s &\models B_a\varphi && \text{iff for all } t \in S : R_a(s, t) \text{ implies } M, t \models \varphi \\ M, s &\models [!\varphi]\psi && \text{iff } M, s \models \varphi \text{ implies } M|\varphi, s \models \psi \end{aligned}$$

where the model restriction $M|\varphi = (S', R', V')$ is defined as $S' = \{s' \in S \mid M, s' \models \varphi\}$ ($= \llbracket \varphi \rrbracket_M$), $R'_a = R_a \cap (S' \times S')$ and $V'(p) = V(p) \cap S'$. \dashv

A complete proof system for this logic for class $\mathcal{S}5$ is presented in [33]. Trivial variations are complete axiomatizations for the model classes \mathcal{K} and $\mathcal{K}45$. The interaction axiom between announcement and belief is:


⁷Unlike in the introduction, p is a Boolean/propositional variable, and not any Boolean proposition.

Definition 4

$$[!\varphi]B_a\psi \leftrightarrow \varphi \rightarrow B_a[!\varphi]\psi \quad \dashv$$


The interaction between announcement and other operators we assume known [44]. It changes predictably in the other logics we present. In the coming sections, we will only vary the dynamic part of the logic, and focus on that completely.

For an example of the semantics of public announcement, consider a situation wherein an agent b is uncertain about p , and receives the information that p . The initial uncertainty requires a model consisting of two states, one where p is true and one where p is false. In view of the continuation, we draw all accessibility relations. For convenience, a state has been given the value of the atom true there as its name. The actual state is underlined.



In the actual state p is true, and from the actual state two states are accessible: that p -state and the $\neg p$ -state. Therefore, the agent does not believe p (as there is an accessible $\neg p$ -state) and does not believe $\neg p$ either (as there is an accessible p -state). She is uncertain whether p . The announcement $!p$ results in a restriction of the epistemic state to the p -state (it is false in the other state). On the right, the agent believes that p . In the initial epistemic state it is therefore true that $p \wedge \neg(B_b p \vee B_b \neg p) \wedge [!p]B_b p$. Note that both on the left and on the right the accessibility relation is an equivalence relation. Indeed, the class $\mathcal{S5}$ is closed under truthful public announcements. The example models change of knowledge rather than the weaker change of belief.

The class $\mathcal{KD45}$ is not closed under truthful public announcements. Let us consider another example. Suppose agent a incorrectly believes p and wishes to process the truthful public announcement that $\neg p$.



On the left, we have that $\neg p \wedge B_a p$ is true. The new information $!\neg p$ results in eliminating the p -state, and consequently agent a 's accessibility relation becomes empty. She believes everything. The model on the left is $\mathcal{KD45}$, but the model on the right is not $\mathcal{KD45}$, it is not serial. However, the class $\mathcal{K45}$ is closed under truthful public announcements.

3 Logic of truthful and lying public announcements

The logic of lying public announcements complements the logic of truthful public announcements. They are inseparably tied to one another. For a clear link we need to use an alternative semantics for truthful public announcement logic. The results in this section are mainly from [45].

We expand the language of truthful public announcement logic with another inductive construct $[i\varphi]\psi$, for ‘after lying public announcement of φ , formula ψ (is true)’; in short ‘after the lie that φ , ψ ’. This is the language $\mathcal{L}(!, i)$.

Definition 5 (Language)

$$\mathcal{L}(!, i) \ni \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \psi) \mid B_a\varphi \mid [!\varphi]\psi \mid [i\varphi]\psi \quad \dashv$$

Truthful public announcement logic is the logic to model the revelations of a benevolent god, taken as the truth without questioning. The announcing agent is not modelled in public announcement logic, but only the effect of its announcements on the audience, the set of all agents. Consider a *false* public announcement, made by a malevolent entity, the devil. Everything he says is false. Everything is a lie. As in religion, god and the devil are inseparable and should be modelled simultaneously.

As a semantics for this logic we employ an alternative to the semantics for public announcement logic from the previous section. This alternative is the semantics of conscious updates [17].⁸ When announcing φ , instead of eliminating states where φ does not hold, one eliminates for each agent those pairs of its accessibility relation where φ does not hold in the second argument of the pair. The effect of the announcement of φ is that only states where φ is true are accessible for the agents. In other words, the semantics is arrow eliminating instead of state eliminating. We call this the *believed public announcement*. New information is accepted by the agents independent from the truth of that information.

Definition 6 (Semantics of believed public announcement)

$$M, s \models [believed\ public\ announcement\ that\ \varphi]\psi \quad \text{iff} \quad M^\varphi, s \models \psi$$

where epistemic model M^φ is as M except that (with S the domain of M)

$$R_a^\varphi := R_a \cap (S \times \llbracket \varphi \rrbracket_M). \quad \dashv$$

In [45], the believed public announcement of φ is called manipulative update with φ . The original proposal there is to view believed public announcement of φ as non-deterministic choice (as in action logics and PDL-style logics) between truthful public announcement of φ and lying public announcement of φ .

Definition 7 (Semantics of truthful and lying public announcement)

$$\begin{aligned} M, s \models [!\varphi]\psi & \quad \text{iff} \quad M, s \models \varphi \text{ implies } M^\varphi, s \models \psi \\ M, s \models [i\varphi]\psi & \quad \text{iff} \quad M, s \models \neg\varphi \text{ implies } M^\varphi, s \models \psi \end{aligned} \quad \dashv$$

⁸A public announcement of φ is a conscious update with the test $?\varphi$ for the entire set of agents A (such updates can be for any subset of agents). In [17] these updates are interpreted on non-wellfounded sets, namely on rooted infinite trees, such as the tree unwinding of a pointed $S5$ model. We present the simpler semantics for conscious update of [23]. We note that [17] and [33] were independent proposals for the semantics of public announcements.

If we now define $[\!|\varphi|\!] \psi$ by abbreviation as $[\!|\varphi|\!] \psi \wedge [\!|\varphi|\!] \psi$, then $[\!|\varphi|\!]$ has the semantics of ‘believed public announcement that φ ’. The non-determinism of this operator (it has two executions, one when φ is true and one when φ is false) comes to the fore if we write it in the dual form: $\langle [\!|\varphi|\!] \rangle \psi \leftrightarrow \langle [\!|\varphi|\!] \rangle \psi \vee \langle [\!|\varphi|\!] \rangle \psi$; in other words, we can view $[\!|\varphi|\!]$ as some $[\!|\varphi|\!] \cup [\!|\varphi|\!]$, where \cup is non-deterministic choice.

The following result justifies that it is not ambiguous to write $[\!|\varphi|\!]$ for ‘arrow elimination’ truthful public announcement but also for ‘state elimination’ truthful public announcement.

Proposition 8 ([23, 45]) Let $s \in \mathcal{D}(M)$ and $M, s \models \varphi$. Then

$$(M^\varphi, s) \Leftrightarrow (M|_\varphi, s). \quad \dashv$$

The symbol \Leftrightarrow stands for ‘is bisimilar to’, a well-known notion that guarantees that the models cannot be distinguished in the logical language [10].

State elimination seems simpler than arrow elimination. Having gone through the trouble of reinterpreting truthful public announcement in arrow elimination semantics, why did we not proceed in the other direction and reinterpret lying public announcement in state elimination semantics? This is not possible! Consider a model wherein all belief is correct, i.e., a model with equivalence relations encoding the beliefs of the agents. A state elimination semantics preserves equivalence, whereas lying inevitably introduces false beliefs: states not accessible to themselves.

The axiom for belief after truthful public announcement remains what it was (Definition 4, and using Proposition 8) and the axiom for the reduction of belief after lying is as follows.

Definition 9 (Axiom for belief after lying [45])

$$[\!|\varphi|\!] B_a \psi \leftrightarrow \neg \varphi \rightarrow B_a [\!|\varphi|\!] \psi. \quad \dashv$$

After the lying public announcement that φ , agent a believes that ψ , if and only if, on condition that φ is false, agent a believes that ψ after truthful public announcement that φ . To the credulous person who believes the lie, the lie appears to be the truth.

Proposition 10 ([45]) The axiomatization of the logic of truthful and lying public announcements is complete (for the model class \mathcal{K} and for the model class $\mathcal{K}45$). \dashv

Proof As for the logic of truthful public announcements, completeness is shown by a reduction argument. All formulas in $\mathcal{L}([\!|\cdot|\!], [\!|\cdot|\!])$ are equivalent to formulas in \mathcal{L} (epistemic logic). By means of equivalences such as in the axiom for belief after lying one can rewrite each formula to an equivalent one without announcement operators. (As the class $\mathcal{K}45$ is not closed under updates with announcements, the logic is not complete for that class.) \square

For an example, we show the effect of truthful and lying announcement of p to an agent b in the model with uncertainty about p . The actual state must be different in these models:

when lying that p , p is in fact false, whereas when truthfully announcing that p , p is in fact true. For lying we get

$$b \begin{array}{c} \curvearrowright \\ \neg p \\ \curvearrowleft \end{array} \leftarrow b \rightarrow p \begin{array}{c} \curvearrowleft \\ p \\ \curvearrowright \end{array} b \xRightarrow{!p} \neg p \text{ --- } b \rightarrow p \begin{array}{c} \curvearrowleft \\ p \\ \curvearrowright \end{array} b$$

whereas for truthtelling we get

$$b \begin{array}{c} \curvearrowright \\ \neg p \\ \curvearrowleft \end{array} \leftarrow b \rightarrow \underline{p} \begin{array}{c} \curvearrowleft \\ p \\ \curvearrowright \end{array} b \xRightarrow{!p} \neg p \text{ --- } b \rightarrow \underline{p} \begin{array}{c} \curvearrowleft \\ p \\ \curvearrowright \end{array} b$$

The reduction principle in [17, 23] for the interaction between belief and believed announcement is, in terms of our language, $[\![\varphi]\!]B_a\varphi \leftrightarrow B_a(\varphi \rightarrow [\![\varphi]\!]\psi)$. This seems to have a different shape, as the modal operator binds the entire implication. But it is indeed valid in our semantics (a technical detail we did not find elsewhere).

Proposition 11 $\models [\![\varphi]\!]B_a\varphi \leftrightarrow B_a(\varphi \rightarrow [\![\varphi]\!]\psi)$ ⊢

Proof

$$\begin{aligned} [\![\varphi]\!]B_a\psi &\Leftrightarrow [\![\varphi]\!]B_a\psi \wedge [\![\varphi]\!]B_a\psi \\ &\Leftrightarrow^\# (\varphi \rightarrow B_a[\![\varphi]\!]\psi) \wedge (\neg\varphi \rightarrow B_a[\![\varphi]\!]\psi) \\ &\Leftrightarrow B_a[\![\varphi]\!]\psi \\ &\Leftrightarrow B_a(\varphi \rightarrow [\![\varphi]\!]\psi) \\ &\Leftrightarrow^* B_a((\varphi \rightarrow [\![\varphi]\!]\psi) \wedge (\varphi \rightarrow [\![\varphi]\!]\psi)) \\ &\Leftrightarrow B_a(\varphi \rightarrow ([\![\varphi]\!]\psi \wedge [\![\varphi]\!]\psi)) \\ &\Leftrightarrow B_a(\varphi \rightarrow [\![\varphi]\!]\psi) \end{aligned}$$

The $\#$ -ed equivalence holds because from the semantics of truthful and lying announcement directly follows that $[\![\varphi]\!]\psi \leftrightarrow (\varphi \rightarrow [\![\varphi]\!]\psi)$ and $[\![\varphi]\!]\psi \leftrightarrow (\neg\varphi \rightarrow [\![\varphi]\!]\psi)$. The $*$ -ed equivalence holds because

$$\begin{aligned} \varphi \rightarrow [\![\varphi]\!]\psi &\Leftrightarrow \varphi \rightarrow (\neg\varphi \rightarrow [\![\varphi]\!]\psi) \\ &\Leftrightarrow (\varphi \wedge \neg\varphi) \rightarrow [\![\varphi]\!]\psi \\ &\Leftrightarrow \perp \rightarrow [\![\varphi]\!]\psi \\ &\Leftrightarrow \top \end{aligned}$$


□

The logic of truthful and lying public announcement satisfies the property ‘substitution of equivalents’, which we have used repeatedly in the proof of Proposition 11, but the logic does not satisfy the property ‘substitution of variables’ (substitution of formulas for propositional variables preserves validity). For example, $[\![p]\!]p$ is valid but, clearly, $[\![p \wedge \neg B_ap]\!](p \wedge \neg B_ap)$ is invalid.


4 Agent announcement logic

In the logic of lying and truthful public announcements, the announcing agent is an outside observer and is implicit. Therefore, it is also implicit that it believes that the announcement is false or true. In multi-agent epistemic logic, it is common to formalize ‘agent a truthfully announces φ ’ as ‘the outside observer truthfully announces $B_a\varphi$ ’. However, ‘agent a lies that φ ’ cannot be modelled as ‘the outside observer lies that $B_a\varphi$ ’. This is the main reason for a logic of lying.


For a counterexample, consider an epistemic state with equivalence relations, modelling knowledge, where b does not know whether p , a knows whether p , and p is true. Agent a is in the position to tell b the truth about p . A truthful public announcement of B_ap (arrow elimination semantics, Definition 7) indeed simulates that a truthfully and publicly announces p .



Given the same model, now suppose p is false, and that a lies that p . A lying public announcement of B_ap (it satisfies the required precondition $\neg B_ap$) does not result in the desired information state, because this makes agent a believe her own lie. In fact, as she already knew $\neg p$, this makes a ’s beliefs inconsistent.



Instead, a lie by a to b that p should have the following effect:



After this lie we have that a still believes that $\neg p$, but that b believes that p . (We even have that b believes that b and a have common belief of p , see Section 9.) We satisfied the requirements of a lying agent announcement, for believed lies.

The precondition for agent a truthtelling that φ is $B_a\varphi$ and the precondition for agent a lying that p is $B_a\neg\varphi$. Another form of announcement is *bluffing*. You are bluffing that φ , if you say that φ but are uncertain whether φ . The precondition for agent a bluffing is therefore $\neg(B_a\varphi \vee B_a\neg\varphi)$. If belief is implicit, we had only two preconditions for announcing φ : φ and $\neg\varphi$, for truthtelling and lying. The ‘third’ would be $\neg(\varphi \vee \neg\varphi)$ which is \perp . The devil can lie, but it cannot bluff.

The postconditions of these three types of announcement are a function of their effect on the accessibility relation of the agents. The effect is the same for all three types. However, it is different for the speaker and for the addressee(s).

- States where φ was believed by speaker a remain accessible to speaker a ;

- States where φ was not believed by speaker a remain accessible to speaker a ;
- States where φ was believed by speaker a remain accessible to addressee b ;
- States where φ was not believed by speaker a are no longer accessible to addressee b .

These requirements are embodied by the following syntax and semantics.

Definition 12 The *language of agent announcement logic* is defined as

$$\mathcal{L}(!_a, i_a, i!_a) \ni \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \psi) \mid B_a\varphi \mid [!_a\varphi]\psi \mid [i_a\varphi]\psi \mid [i!_a\varphi]\psi \quad \dashv$$

The inductive constructs $[!_a\varphi]\psi$, $[i_a\varphi]\psi$, and $[i!_a\varphi]\psi$ stand for, respectively, a truthfully announces φ , a is lying that φ , and a is bluffing that φ ; where agent a addresses all other agents b .

Definition 13

$$\begin{aligned} M, s \models [!_a\varphi]\psi & \text{ iff } M, s \models B_a\varphi \text{ implies } M_a^\varphi, s \models \psi \\ M, s \models [i_a\varphi]\psi & \text{ iff } M, s \models B_a\neg\varphi \text{ implies } M_a^\varphi, s \models \psi \\ M, s \models [i!_a\varphi]\psi & \text{ iff } M, s \models \neg(B_a\varphi \vee B_a\neg\varphi) \text{ implies } M_a^\varphi, s \models \psi \end{aligned}$$

where M_a^φ is as M except for the accessibility relation R' defined as (S is the domain of M , and $a \neq b$)

$$\begin{aligned} R'_a &:= R_a \\ R'_b &:= R_b \cap (S \times \llbracket B_a\varphi \rrbracket_M). \end{aligned} \quad \dashv$$

The principles for a truthtelling, lying, or bluffing to b are as follows. The belief consequences for the speaker are different from the belief consequences for the addressee(s).

Definition 14 (Axioms for the belief consequences of announcements)

$$\begin{aligned} [!_a\varphi]B_b\psi & \leftrightarrow B_a\varphi \rightarrow B_b[!_a\varphi]\psi \\ [!_a\varphi]B_a\psi & \leftrightarrow B_a\varphi \rightarrow B_a[!_a\varphi]\psi \\ [i_a\varphi]B_b\psi & \leftrightarrow B_a\neg\varphi \rightarrow B_b[i_a\varphi]\psi \\ [i_a\varphi]B_a\psi & \leftrightarrow B_a\neg\varphi \rightarrow B_a[i_a\varphi]\psi \\ [i!_a\varphi]B_b\psi & \leftrightarrow \neg(B_a\varphi \vee B_a\neg\varphi) \rightarrow B_b[i!_a\varphi]\psi \\ [i!_a\varphi]B_a\psi & \leftrightarrow \neg(B_a\varphi \vee B_a\neg\varphi) \rightarrow B_a[i!_a\varphi]\psi \end{aligned}$$


In other words, the liar knows that he is lying, but the dupe he is lying to, believes that the liar is telling the truth. The principles for truthtelling and bluffing are similar to that for lying, but with the obvious different conditions on the right hand side of the equivalences. The bluffer knows that he is bluffing, but the dupe he is bluffing to, believes that the bluffer is telling the truth. And in case the announcing agent is truthful, there is no discrepancy, both the speaker and the addressee believe that the consequences are those of truthtelling.

Proposition 15 The axiomatization of the logic of agent announcements is complete. \dashv

Proof Just as in the previous logics (see Proposition 10) completeness is shown by a reduction argument. All formulas in $\mathcal{L}(!_a, i_a, i!_a)$ are equivalent to formulas in epistemic logic. In the axioms above, the announcement operator is (on the right) always pushed further down into any given formula. As before, the result holds for model classes \mathcal{K} , $\mathcal{K}45$ and $\mathcal{S}5$, but not for $\mathcal{KD}45$.

An alternative, indirect, completeness proof is that agent announcement logic is action model logic for a given action model, as will be explained in Section 6. \square

As an example illustrating the difference between a truthtelling, lying and bluffing agent announcement we present the following model wherein the addressee b , hearing the announcement of p by agent a , considers all three possible. In fact, a is bluffing, and a 's announcement that p is false. After the announcement, b incorrectly believes that p , but a is still uncertain whether p . If the bottom-left state had been the actual state, a would have been lying that p , and if the bottom-right state had been the actual state, it would have been a truthful announcement by a that p .



Unbelievable announcements The axiomatization of the logic of agent announcements is incomplete for $\mathcal{KD}45$ (see the proof of Proposition 15) because of the problem of unbelievable announcements. In Sections 7 and 8 we present alternative logics wherein believable announcements (announcements of φ to addressee b such that $\neg B_b \neg \varphi$ is true) are treated differently from unbelievable announcements (such that $B_b \neg \varphi$ is true). These logics are complete for class $\mathcal{KD}45$.

Outside observer Consider a depiction of an epistemic model. The outside observer is the guy or girl looking at the picture: you, the reader. She can see all different states. She has no uncertainty and her beliefs are correct. It is therefore that her truthful announcements are true and that her lying announcements are false. It is also therefore that ‘truthful public announcement logic’ is not a misnomer, it is indeed the logic of how to process new information that is true.

We can model the outside observer as an agent gd , for ‘god or the devil’. (The model does not need to be a bisimulation contraction.)

Proposition 16 Given an epistemic model M , let $gd \in A$ be an agent with an accessibility relation that is the identity on M . Then $M_{gd}^\varphi = M^\varphi$. Let φ, ψ not contain announcement operators, then $[\![_{gd}\varphi]\!\psi$ is equivalent to $[\![_\varphi]\!\psi$, and $[\![_{gd}\varphi]\!\psi$ is equivalent to $[\![_\varphi]\!\psi$. \dashv

Proof In Definition 13, the accessibility of the addressees is adjusted to $R'_b := R_b \cap (S \times \llbracket B_{gd}\varphi \rrbracket_M)$. As R_{gd} is the identity, $B_{gd}\varphi$ is equivalent to φ . So, $R'_b := R_b \cap (S \times \llbracket \varphi \rrbracket_M)$, as in Definition 6. \square

Lying about beliefs Agents may announce factual propositions (Boolean formulas) but also modal propositions, and thus be lying and bluffing about them. In the consecutive numbers riddle the announcements ‘I know your number’ and ‘I do not know your number’ are modal propositions, and the agents may be lying about those.⁹ (Details in Section 5.)

For our target agents, that satisfy introspection (so that $B_a B_a \varphi \leftrightarrow B_a \varphi$ and $B_a \neg B_a \varphi \leftrightarrow \neg B_a \varphi$ are validities), the distinction between bluffing and lying seems to become blurred. If I am uncertain whether p , I would be bluffing if I told you that p , but I would be lying if I told you that I believe that p . The announcement that p satisfies the precondition $\neg(B_a p \vee B_a \neg p)$. It is bluffing that p (it is $!\!_a p$). But the announcement that $B_a p$ satisfies the precondition $B_a \neg B_a p$, the negation of the announcement. It is lying that $B_a p$ (it is $!\!_a B_a p$). (Proof: $\neg(B_a p \vee B_a \neg p)$ is equivalent to $\neg B_a p \wedge \neg B_a \neg p$, from which follows $\neg B_a p$, and with negative introspection $B_a \neg B_a p$.) We would prefer to call both bluffing, and ‘ a announces that $B_a p$ ’ strictly or really ‘ a announces that p ’. A general solution to avoid such ambiguity involves more than merely stripping a formula of an outer B_a operator: a announcing that $B_a B_a p$ should also strictly be a announcing that p , and a announcing that $B_a p \wedge B_a q$ should strictly be a announcing that that $p \wedge q$. We need recursion. In the following definition and proposition we assume an agent with consistent beliefs.

Definition 17 (Strictly lying) An announcement by a that φ is *strict* iff φ is equivalent to \top or φ is equivalent to $\bigvee_i \psi_i^a \wedge \bigwedge_i \neg B_a \neg \psi_i^a$ where all ψ_i^a are alternating disjunctive forms (see proof below) for agents other than a . An agent announcement $!\!_a \varphi$ is *strictly lying* iff there is a φ' equivalent to φ such that φ' is strict and $!\!_a \varphi'$ is a lying agent announcement. (Similarly for *strictly bluffing*.) \dashv

Proposition 18 For each $\varphi \in \mathcal{L}(!_a, !_a, !_a)$ there is an equivalent $\psi \in \mathcal{L}$ that is strict. \dashv

⁹In social interaction, untruthfully announcing *epistemic* propositions is not always considered lying with the negative moral connotation. Suppose we work in the same department and one of our colleagues, X , is having a divorce. I know this. I also know that you know this. But we have not discussed the matter between us. I can bring up the matter in conversation by saying ‘You know that X is having a divorce!’. But this is unwise. You may not be willing to admit your knowledge, because X ’s husband is your friend, which I have no reason to know; etc. A better strategy for me is to say ‘You may not know that X is having a divorce’. This is a lie. I do not consider it possible that you do not know that. But, unless we are very good friends, you will not laugh in my face to that and respond with ‘Liar!’. Could it be that lies about facts are typically considered worse than lies about epistemic propositions, and that the more modalities you stack in your lying announcement, the more innocent the lie becomes?

Proof We first define the *alternating disjunctive form* (*adf*). An *adf* is a disjunction of formulas of shape $\psi_0 \wedge \bigwedge_{b \in A'} (B_b \bigvee_i \psi_i^b \wedge \bigwedge_i \neg B_b \neg \psi_i^b)$, where $\psi_0 \in \mathcal{L}$, $A' \subseteq A$, and where ψ_i^b is an *adf* for a group $A'' \subset A$ of agents other than b (the ‘alternating’ in alternating disjunctive forms). The indices i range over $1, \dots, n$ ($n \geq 1$). We may assume that for each b some ψ_i^b is non-trivial (i.e., not equivalent to \top), because otherwise we can take $A' \setminus b$ in the conjunction above. Each formula in multi-agent $\mathcal{KD45}$ is equivalent to an *adf* [20].¹⁰

Now for the proof. Let $\varphi \in \mathcal{L}(!_a, !_a, !_a)$. Determine an equivalent $\varphi' \in \mathcal{L}$ (i.e., rewrite φ into a multi-agent epistemic formula, without announcement operators). Let ψ be an *adf* equivalent to $B_a \varphi'$. It is elementary to see that ψ must be either equivalent to \top or have the form $B_a \bigvee_i \psi_i^a \wedge \bigwedge_i \neg B_a \neg \psi_i^a$ and where some ψ_i^a is non-trivial. We now observe that

$$\begin{aligned} \psi &\Leftrightarrow B_a \bigvee_i \psi_i^a \wedge \bigwedge_i \neg B_a \neg \psi_i^a \\ &\Leftrightarrow B_a \bigvee_i \psi_i^a \wedge \bigwedge_i B_a \neg B_a \neg \psi_i^a \\ &\Leftrightarrow B_a (\bigvee_i \psi_i^a \wedge \bigwedge_i \neg B_a \neg \psi_i^a) \end{aligned}$$

and that $\bigvee_i \psi_i^a \wedge \bigwedge_i \neg B_a \neg \psi_i^a$ is strict. \square

Mistakes and lies There are two sorts of mistaken beliefs.

The first kind is when I believe $p \rightarrow \neg q$ and I believe p , and, when questioned about my belief in q , confidently state that I believe that q . This is a mistaken belief. The formula $B_a q$ should be false, because $B_a \neg q$ is a deductive consequence of $B_a(p \rightarrow \neg q)$ and $B_a p$. To explain this sort of mistaken belief we need to resort to bounded rationality. We do not do that.

The second kind of mistaken belief is when I believe that p but when in fact p is false. Fully rational agents can make such mistakes. This, we can address. It amounts to the truth of $\neg p \wedge B_a p$.

What is the difference between a lie and a mistake? I am lying if I say φ and believe $\neg \varphi$ (independently from the truth of φ), whereas I am mistaken if I say φ and believe φ , but φ is false. Let a be the speaker, then the precondition of lying that φ is $B_a \neg \varphi$ and the precondition of a mistaken truthful announcement that φ is $\neg \varphi \wedge B_a \varphi$. The speaker can therefore distinguish a lie from a mistake.

How about the addressee? Clearly, a $\mathcal{KD45}$ agent cannot distinguish another agent lying from being mistaken, because it can itself be mistaken about the announced proposition and about the speaker’s beliefs of that proposition. We can then only observe that if a says φ given that $B_b(\neg \varphi \wedge B_a \varphi)$, then b believes a to be mistaken, whereas if a says φ given that $B_b B_a \neg \varphi$, then b believes a to be lying.

But if both knowledge and belief play a role then the addressee can distinguish a lie from a mistake. A standard assumption in many games (therefore called ‘fair games’)

¹⁰Reported in [20], an *adf* is a multi-agent generalization of the (single-agent) disjunctive form as in [30, p.35] (where it is called *S5 normal form*), and a special case of the disjunctive form as in [14]. An *adf* contains no stacks of B_a operators without an intermediary B_b operator for another agent, i.e., if $B_a \chi$ is a subformula of an *adf* and $B_a \chi''$ is a subformula of χ then there is an agent $b \neq a$ and a χ' such that $B_b \chi'$ is a subformula of χ and $B_a \chi''$ is a subformula of χ' .

and in many other distributed systems, is initial common knowledge of all agents of the uncertainty about the system. (I *know* — I can assume to know — that you do not know my card. If not, it's not fair, it was cheating.) Part of that initial common knowledge is of facts or otherwise positive formulas, persisting after every update. In such a system, beliefs that are not knowledge can only be induced by incorrect updates (lies). For such an addressee, a mistake is that a says φ when $K_b(\neg\varphi \wedge B_a\varphi)$, whereas a lie is that a says φ when $K_bK_a\neg\varphi$. The consecutive numbers riddle gives an example of such a lie and such a mistake.

5 Lying in the consecutive numbers riddle

In this section we give an extended example of agent announcement logic. We recall the consecutive numbers riddle.

Anne and Bill are each going to be told a natural number. Their numbers will be one apart. The numbers are now being whispered in their respective ears. They are aware of this scenario. Suppose Anne is told 2 and Bill is told 3.

The following truthful conversation between Anne and Bill now takes place:

- *Anne: “I do not know your number.”*
- *Bill: “I do not know your number.”*
- *Anne: “I know your number.”*
- *Bill: “I know your number.”*

Explain why is this possible.


We first analyze the riddle in public announcement logic with the state elimination semantics. Although the riddle goes back to 1953 [25], this analysis dates from 1984 [46] and played an important role in the development of the area of dynamic epistemic logic.¹¹ The initial model encodes that the numbers are consecutive and consists of two disconnected countable parts, one where the agents have common knowledge (correct common belief) that a 's number is odd and b 's number is even, and another one where the agents have common knowledge that a 's number is even and b 's number is odd. As before, the actual state is underlined. This is the state $(2, 3)$, where a is told 2 and b is told 3. In this simplified visualization we use the convention that symmetry and reflexivity are assumed.

$(1,0) - a - (1,2) - b - (3,2) - a - (3,4) \dots$ $(0,1) - b - (2,1) - a - \underline{(2,3)} - b - (4,3) \dots$


¹¹In [46] the consecutive numbers riddle is called the Conway Paradox, after [13]. That name that has since stuck. Conway (and Paterson) indeed present an epistemic riddle in [13], but not the consecutive numbers riddle. This was a mistaken attribution, as Peter van Emde Boas (an author of [46]) told us.

To determine the model restriction for an announcement, given a number pair (m, n) let m_a stand for ‘Anne is told number m ’ and let n_b stand for ‘Bill is told number n ’. Each state corresponds to a number pair (m, n) . To process Anne’s first announcement we determine if $\neg B_a n_b$ in that state. If so, the state is kept. Otherwise, the state is eliminated. For Bill’s subsequent announcement we do this for $\neg B_b m_a$. Etc.: the third time for $B_a n_b$, and the last time for $B_b m_a$.¹² We successively process all four announcements. The epistemic state resulting from an announcement is shown immediately after it.


- Anne: “I do not know your number.”



- Bill: “I do not know your number.”



- Anne: “I know your number.”




- Bill: “I know your number.”

The last announcement does not make a difference anymore, as it is already common knowledge that Anne and Bill know each other’s number.


Consecutive numbers with lying Next, we show two different scenarios for the consecutive numbers riddle with lying. This is agent truthtelling and agent lying, the epistemic actions defined as $!_a \varphi$ and $!_a \neg \varphi$ (Definition 12). With lying, the riddle involves (the speaker) feigning knowledge and (the addressee) incorrectly believing something to be knowledge, so we move from knowledge to belief. In the communication only occur ‘I know your number’ and ‘I do not know your number’. Bluffing can therefore not be demonstrated: introspective agents believe their uncertainty and believe their beliefs. As we are reasoning from the actual state $(2, 3)$, to simplify matters we do not depict the disconnected upper part of the model. And as beliefs may be incorrect, we show all arrows.

¹²Although the riddle is a standard illustration of truthful public announcement logic, these announcements cannot actually be formalized in the language of that logic. For the first announcement it would be the infinitary expression $\bigwedge_{n \in \mathbb{N}} (n_b \rightarrow \neg B_a n_b)$ (the expression holds for all $n \neq 1$). This is not a formula in $\mathcal{L}(!)$. To demonstrate a and b lying to each other this does not matter.

First scenario The first scenario consists of Anne lying in her first announcement. Bill does not believe Anne's announcement: his accessibility relation has become empty. Bill's beliefs are therefore no longer consistent. Here, the analysis stops. We do not treat Bill's 'announcement' "That's a lie" as a permitted move in the game. On the assumption of initial common knowledge Bill *knows* that Anne was lying and not mistaken.




- Anne: "I know your number." *Anne is lying*




- Bill: "That's a lie."

Second scenario In the second scenario Anne initially tells the truth, after which Bill is lying, resulting in Anne mistakenly concluding (and announcing) that she knows Bill's number: she believes it to be 1. This mistaken announcement by Anne is informative to Bill. He learns from it (correctly) that Anne's number is 2, something he didn't know before.


In the second scenario, Bill gets away with lying, because Anne considered it possible that he told the truth. Bill knows (believes correctly, and justifiably) that Anne's second announcement was a mistake and not a lie. The justification is, again, common knowledge in the initial epistemic state.




- Anne: "I do not know your number."



- Bill: "I know your number." *Bill is lying*




- Anne: "I know your number." *Anne is mistaken.*




Once the agents consider it possible that the other is lying, every announcement can be either truthful or lying. This is obvious, because (just as for the semantics of ‘believed public announcement’) the structural transformation induced by agent announcement depends on the model only, and not the actual state. The two scenarios can therefore be executed in any given state of the epistemic model. Whether the announcements are lying or truthtelling depends on the actual state (only).

For example, in the first scenario, if the initial state had been $(4, 3)$, it would also have been a lie, and Bill would have similarly detected the lie. If the initial state had been $(2, 1)$, it would still have been a lie, but Bill would not have known that. And if the initial state had been $(0, 1)$, it would have been the truth. Similar variations can be outlined for the second scenario.

Instead of considering Bill’s “That’s a lie” as terminating the game by observing an illegal move (breach of protocol), we could consider modelling this as an informative (agent) announcement. It amounts to Bill announcing \perp , i.e., as the truthful agent announcement $!B_b\perp$ (‘I believe a false statement’). In $(2, 3)$, the formula $B_b\perp$ is true. In a state where Bill’s beliefs are consistent, it is false. The informative consequences are as follows.



- Bill: “That’s a lie.”



Of course, this way we declare it a permitted move in the consecutive numbers game, and Bill could have lied in his announcement! Consider actual number pair $(0, 1)$. Anne truthfully announces “I know your number.” Following Bill’s response “That’s a lie”, Anne will now even more indignantly contest: “That cannot be, *you* are lying!”, i.e., “Your announcement “That’s a lie” was a lie.” We recall that given initial common knowledge Anne knows that Bill is in fact lying. There is no uncertainty. Lying games may benefit from analysis in game theory, by associating a high negative payoff to the discovery of lying, see Section 9.

6 Action models and lying

Whether I am truthfully announcing φ to you, or am lying that φ , or am bluffing that φ , to you it all appears as the same announcement $\varphi!$. Whereas to me, the speaker, they all

appear to be different announcements: I know whether I am truthful, lying, or bluffing. Different agents have different perspectives on this action. Action models [7] are a familiar way to formalize uncertainty about actions in that form of different perspectives on ‘what the real action is’. The action model for truthful public announcement (state elimination semantics) can be viewed as a singleton action model. This is well-known. We can view truthful and lying public announcement (‘believed announcement’, the arrow elimination semantics) as the different points, respectively, of a two-point action model. This is somewhat less well-known. (The modelling of believed announcements [17] as an action model was suggested in [39, 23].) We can also view truthful, lying and bluffing agent announcement as the respective different points of a three-point action model. These should be seen as alternative descriptions of the logics of lying in terms of a well-known framework. It has the additional advantage of independently validating the various axioms for lying and bluffing, and providing alternative completeness proofs (action model logic is complete). All the following definitions can be found in any standard introduction to action models, such as [44].

An action model is a structure like a Kripke model but with a precondition function instead of a valuation function.

Definition 19 (Action model) An *action model* $M = (S, R, \text{pre})$ consists of a *domain* S of *actions*, an *accessibility function* $R : A \rightarrow \mathcal{P}(S \times S)$, where each R_a is an accessibility relation, and a *precondition function* $\text{pre} : S \rightarrow \mathcal{L}_X$, where \mathcal{L}_X is a logical language. A pointed action model is an *epistemic action*. \dashv

A truthful public announcement of φ (state elimination semantics) is a singleton action model with precondition φ and with the single action accessible to all agents.


Performing an epistemic action in an epistemic state means computing their restricted modal product. This product encodes the new state of information.

Definition 20 (Update of an epistemic state with an action model) Given an epistemic state (M, s) where $M = (S, R, V)$ and an epistemic action (M, s) where $M = (S, R, \text{pre})$. Let $M, s \models \text{pre}(s)$. The update $(M \otimes M, (s, s))$ is an epistemic state where $M \otimes M = (S', R', V')$ such that

$$\begin{aligned} S' &= \{(t, \mathbf{t}) \mid M, t \models \text{pre}(\mathbf{t})\} \\ ((t, \mathbf{t}), (t', \mathbf{t}')) \in R'_a &\text{ iff } (t, t') \in R_a \text{ and } (\mathbf{t}, \mathbf{t}') \in R_a \\ (t, \mathbf{t}) \in V'(p) &\text{ iff } t \in V(p) \end{aligned} \quad \dashv$$

The domain of $M \otimes M$ is the product of the domains of M and M , but restricted to state/action pairs (t, \mathbf{t}) such that $M, t \models \text{pre}(\mathbf{t})$, i.e., such that the action can be executed in that state. (Note that a state in the resulting domain is no longer an abstract object, as before, but such a state/action-pair.) An agent considers a pair (t, \mathbf{t}) possible in the next information state if she considered the previous state t possible, and the execution of action \mathbf{t} in that state. And the valuations do not change after action execution.

Definition 21 (Action model for truthful and lying pub. ann.) The action model M' for truthful and lying public announcement that φ , where $\varphi \in \mathcal{L}(!, i)$, consists of two actions (suggestively) named $!$ and i , with $\text{pre}(!) = \varphi$ and $\text{pre}(i) = \neg\varphi$, and such that for all agents only action $!$ is accessible. Truthful public announcement of φ is the epistemic action $(M', !)$. Lying public announcement of φ is the epistemic action (M', i) . The action model can be depicted as follows. Preconditions are shown below the actions.



⊢

This terminology is not ambiguous in view of our earlier results, as we have the following.


Proposition 22 ([39, 23])

$$\begin{aligned} M, s \models [! \varphi] \psi & \text{ iff } M \otimes M', (s, !) \models \psi \\ M, s \models [i \varphi] \psi & \text{ iff } M \otimes M', (s, i) \models \psi . \end{aligned}$$

⊢

Proof Elementary, by induction on ψ . □

For an example, we show the execution of (M', i) for ‘the outside observer is lying to a that p ’, in the epistemic state where the agent a is uncertain whether p but where p is false.



Consider this epistemic model (S, R, V) . As before, the states in the epistemic model are named after their valuation, so that $S = \{\neg p, p\}$; the state named $\neg p$ is simply the state where p is false, etc. The modal product (S', R', V') consists of two states; $(\neg p, i) \in S'$ because $M, \neg p \models \text{pre}(i)$, as $\text{pre}(i) = \neg p$, and $(p, !) \in S'$ because $M, p \models p$. Then, $((\neg p, i), (p, !)) \in R'_a$ because $(\neg p, p) \in R_a$ (in M) and $(i, !) \in R_a$ (in the action model M'), etc. It is an artifact of the example that the shape of the action model is the shape of the next epistemic state. That is merely a consequence of the fact that the initial epistemic state has the universal accessibility relation for the agent on a domain of all valuations of the atoms occurring in the precondition (here: a domain of two states, for the two valuations of p).

Possibly more elegantly, we can also show the correspondence established in Proposition 22 from the perspective of a different logic. With an epistemic action (M, s) we can associate a dynamic modal operator $[M, s]$ in a logical language where an enumeration of action model frames is a parameter of the inductive language definition, apart from propositional variables P and agents A .

Definition 23 (Language of action model logic)

$$\mathcal{L}(\otimes) \ni \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \psi) \mid B_a\varphi \mid [\mathbf{M}, \mathbf{s}]\psi \quad \dashv$$

The last clause is in fact inductive, if we realize that the preconditions of all actions in \mathbf{M} , including \mathbf{s} , are also of type formula. For example, truthful public announcement logic is an instantiation of that language for the singleton set of actions $\{!\}$, where we view ‘!’ as an operation with two input formulae φ and ψ and that returns as output the announcement formula $[\![\varphi]\!]\psi$.

Definition 24 (Semantics of $[\mathbf{M}, \mathbf{s}]$)

$$M, s \models [\mathbf{M}, \mathbf{s}]\psi \text{ iff } M, s \models \text{pre}(\mathbf{s}) \text{ implies } M \otimes \mathbf{M}, (s, \mathbf{s}) \models \psi \quad \dashv$$

Given the action model for truthful and lying public announcement that φ of Definition 21, but with $\varphi \in \mathcal{L}(\otimes)$, and given an inductively defined translation $tr : \mathcal{L}(\otimes) \rightarrow \mathcal{L}(!, i)$ with only nontrivial clauses $tr([\mathbf{M}', !]\psi) := [!tr(\varphi)]tr(\psi)$ and $tr([\mathbf{M}', i]\psi) := [itr(\varphi)]tr(\psi)$, the correspondence established in Proposition 22 can be viewed from the perspective of the dynamic modal operators for action models (where we suggestively execute the first step of the translation).

$$\begin{aligned} M, s \models [!tr(\varphi)]tr(\psi) & \text{ iff } M, s \models [\mathbf{M}', !]\psi \\ M, s \models [itr(\varphi)]tr(\psi) & \text{ iff } M, s \models [\mathbf{M}', i]\psi \end{aligned}$$

We now proceed with the presentation of action models for agent announcements.

Definition 25 (Action model for agent announcement) The action model \mathbf{M}'' for agent announcement consists of three actions named $i!_a$, $!_a$, and i_a with preconditions $\neg(B_a\varphi \vee B_a\neg\varphi)$, $B_a\varphi$, and $B_a\neg\varphi$, respectively, where $\varphi \in \mathcal{L}(!_a, i_a, i!_a)$. The announcing agent a has identity access on the action model. To the other agents b only action $!_a$ is accessible. Agent a truthfully announcing φ to all other b is the epistemic action $(\mathbf{M}'', !_a)$ — with precondition $B_a\varphi$, therefore — and similarly lying and bluffing are the action models (\mathbf{M}'', i_a) and $(\mathbf{M}'', i!_a)$.

$$\begin{array}{ccccc} \begin{array}{c} \textcircled{a} \\ \downarrow \\ i!_a \end{array} & \xrightarrow{b} & \begin{array}{c} \textcircled{ab} \\ \downarrow \\ !_a \end{array} & \xleftarrow{b} & \begin{array}{c} \textcircled{a} \\ \downarrow \\ i_a \end{array} \\ \neg(B_a\varphi \vee B_a\neg\varphi) & & B_a\varphi & & B_a\neg\varphi \end{array} \quad \dashv$$

Again, we have the desired correspondence (and this can, again, be formulated in action model logic with an inductive translation).

Proposition 26

$$\begin{aligned} M, s \models [!_a\varphi]\psi & \text{ iff } M \otimes \mathbf{M}'', (s, !_a) \models \psi \\ M, s \models [i_a\varphi]\psi & \text{ iff } M \otimes \mathbf{M}'', (s, i_a) \models \psi \\ M, s \models [i!_a\varphi]\psi & \text{ iff } M \otimes \mathbf{M}'', (s, i!_a) \models \psi \end{aligned} \quad \dashv$$

The action model representations validate the axioms for announcement and belief, for all versions shown; and they justify that these axioms form part of complete axiomatizations. These axioms are instantiations of the more general axiom for an epistemic action followed by a belief, in action model logic. This axiom ([7]) is

$$[M, s]B_a\psi \leftrightarrow \text{pre}(s) \rightarrow \bigwedge_{(s,t) \in R_a} B_a[M, t]\psi$$

In other words, an agent believes ψ after a given action, if ψ holds after any action that is for a indistinguishable from it. For example, in the epistemic action (M', i) , with $\text{pre}(i) = \varphi$, for lying public announcement that φ , $!$ is the only accessible action from action i , and we get

$$[M', i]B_a\psi \leftrightarrow \text{pre}(i) \rightarrow B_a[M', !]\psi$$

and therefore (we recall Definition 9)

$$[i\varphi]B_a\psi \leftrightarrow \neg\varphi \rightarrow B_a[!\varphi]\psi$$

In view of the identification discussed in this section, in the following we will continue to call any dynamic modal operator for belief change an epistemic action, and also standardly represent epistemic actions by their corresponding action models.

7 Unbelievable lies and skeptical agents

If I tell you φ and you already believe the opposite, accepting this information will make your beliefs inconsistent. This is not merely a problem for lying (‘unbelievable lies’) but for any form of information update. One way to preserve consistent beliefs is to reject new information if it is inconsistent with your beliefs. (The other way is to accept them, but to remove inconsistent prior beliefs. See Section 8.) Such agents may be called skeptical. In this section we adjust the logics of truthful and lying public announcement and of agent announcement to sceptical agents. This adjustment is elementary. As this topic is alive in the community of dynamic epistemic logicians we incorporate a review of the literature.

Consistency of beliefs is preserved iff seriality is preserved on epistemic models to interpret these beliefs. For the model classes $\mathcal{KD45}$ and $\mathcal{S5}$ that we target, this is the requirement that the class is closed under information update. The perspective on epistemic actions and action models from Section 6 is instructive. The class of $\mathcal{S5}$ epistemic models is closed under update with $\mathcal{S5}$ epistemic actions, such as truthful public announcements, but the class of $\mathcal{KD45}$ models is *not* closed under update with $\mathcal{KD45}$ epistemic actions (see the last paragraph of Section 2). The action models for truthful and lying public announcement, and for agent announcement, are $\mathcal{KD45}$.


Updates that preserve $\mathcal{KD45}$ have been investigated in [37, 4, 24]. Aucher [4] defines a language fragment that makes you go mad (‘crazy formulas’). The idea is then to avoid that. Steiner [37] proposes that the agent does not incorporate the new information if it already believes to the contrary. In that case, nothing happens. Otherwise, access

to states where the information is not believed is eliminated, just as for believed public announcements.¹³ This solution to model unbelievable lies (and unbelievable truths!) is similarly proposed in the elegant [24], where it is called ‘cautious update’ — a suitable term. We will propose to call such agents *skeptical* instead of cautious.

7.1 Logic of truthful and lying public announcements to skeptical agents

We propose a three-point action model for *truthful and lying public announcement to skeptical agents*, with the semantics motivated by [37, 24].

Definition 27 (Public announcement to skeptical agents) The action model \mathbf{N} for truthful and lying public announcement to skeptical agents consists of three actions named $!^{\text{sc}}$, $!^{\text{sc}}$, and (for lack of a better symbol) $!!^{\text{sc}}$, with preconditions and accessibility relations (for all agents a) as follows.



For $(\mathbf{N}, !^{\text{sc}})$, with $\text{pre}(!^{\text{sc}}) = \varphi \wedge \neg B_a \neg \varphi$, we write $!^{\text{sc}}\varphi$; similarly, we write $!^{\text{sc}}\varphi$ for $(\mathbf{N}, !^{\text{sc}})$ and $!!^{\text{sc}}\varphi$ for $(\mathbf{N}, !!^{\text{sc}})$. ⊢

The difference with the action model for truthful and lying public announcement is that the alternatives φ and $\neg\varphi$ now have an additional precondition $\neg B_a \neg \varphi$: the announcement should be believable. In the action model there is a separate, disconnected, case for unbelievable announcements: precondition $B_a \neg \varphi$. For unbelievable announcements it does not matter whether φ is a lie or is the truth. The agent chooses to discard it either way. It is skeptical.

Definition 28 (Axioms) The principles for public announcements to skeptical agents are:

$$\begin{aligned} [!^{\text{sc}}\varphi]B_a\psi &\leftrightarrow (\varphi \wedge \neg B_a \neg \varphi) \rightarrow B_a[!^{\text{sc}}\varphi]\psi \\ [!^{\text{sc}}\varphi]B_a\psi &\leftrightarrow (\neg\varphi \wedge \neg B_a \neg \varphi) \rightarrow B_a[!^{\text{sc}}\varphi]\psi \\ [!!^{\text{sc}}\varphi]B_a\psi &\leftrightarrow B_a \neg \varphi \rightarrow B_a\psi \end{aligned} \quad \dashv$$

¹³Steiner gives an interesting parable for the case where you do not accept new information. Someone is calling you and is telling you something that you don’t want to believe. What do you do? You start shouting through the phone: ‘What did you say? Is there anyone on the other side? The connection is bad!’ And then you hang up, quickly, before the caller can repeat his message. Thus you create common knowledge (by convention) that the message has been received but its content not accepted.

This analysis comes close but is just off, because to the caller this is indistinguishable from the case where the message is believable to the addressee and would have been accepted. In the context of [37] the analysis appears justifiable because the caller is the outside observer who is not modelled. But the outside observer *knows* whether the receiver will find the message believable (Proposition 16), unlike in Steiner’s scenario.

Proposition 29 The axiomatization of the logic for public announcements to skeptical agents is complete. \dashv

Proof Directly, from the embedding in action model logic. \square

In the case of a single agent, the semantics of an unbelievable announcement of φ leaves the structure of the unbelievable part of the model unchanged. This means that agent a does not change its beliefs. An unbelievable public announcement does not make an informative difference for the skeptical agent.


If there are more agents, an announcement can be believable for one agent and unbelievable for another agent. The current action model is no longer appropriate. For example, for two agents there are seven distinct actions, namely all Boolean combinations of φ , $B_a\neg\varphi$, and $B_b\neg\varphi$, minus one: the cases $\varphi \wedge B_a\neg\varphi \wedge B_b\neg\varphi$ and $\neg\varphi \wedge B_a\neg\varphi \wedge B_b\neg\varphi$ can be merged into a single action with precondition $B_a\neg\varphi \wedge B_b\neg\varphi$, because they are indistinguishable for a and b jointly. Such a two-agent skeptical action model is not a refinement of the action model above. Unlike the above, it is connected. In particular, the axiom $[!!^{sc}\varphi]B_a\psi \leftrightarrow B_a\neg\varphi \rightarrow B_a\psi$ is no longer valid, and even when a believes to the contrary, and therefore does not accept φ , its beliefs may still change in another way. For example of such changing $B_a\psi$, if the announcement is that p and I believe that $\neg p$, but I consider it possible that you are uncertain about p , then after the announcement I consider it possible that you now believe p . So for $B_a\psi = B_a\neg B_a B_b p$, this was true before and false after the announcement.

We consider such matters mere variations, and move on.

7.2 Logic of agent announcements to skeptical agents

The analysis becomes more interesting for agent announcements, namely when the speaker is uncertain whether her lie will be believed by the addressee, as in the consecutive numbers riddle. Even if the addressee already believes $\neg p$, he may consider it possible that the speaker is truthfully announcing p and is not lying. The addressee b then merely concludes that the speaker a must be mistaken in her truthful announcement of p , and thus believes $\neg p \wedge B_a p$. If a was lying, this belief is mistaken. The charitable addressee discards the option that the speaker was lying, with precondition $B_a\neg p$.

Definition 30 (Action model for agent announcement to skeptics) The action model N' for agent announcements from speaker a to skeptical addressee(s) b is as follows. Assume transitivity of accessibility relations.



For (N', i_a^{sc}) with precondition $B_a\neg\varphi \wedge \neg B_b\neg\varphi$ we write $i_a^{\text{sc}}\varphi$, etc. (Similarly for public announcement to skeptical agents.) \dashv

This action model encodes that a knows/believes if she's bluffing, lying, or truthful, but is uncertain if her announcement is affecting b 's beliefs: she cannot (from the appearance of the action itself, discarding prior knowledge) distinguish between the actions with preconditions $B_b\neg\varphi$ and $\neg B_b\neg\varphi$. In case agent b already believed $\neg\varphi$, he is indifferent between the three alternatives truthtelling, lying, and bluffing that φ . But in case b considered φ possible, as before, he believes a to be truthful about φ .


We only give the reduction axioms for lying, for this logic, and not the remaining axioms. The more interesting lying axiom is the second one. It formalizes that the lying agent a may be uncertain if the lie was believable to the addressee b (the two cases on the right), even if in fact (on the left) the lie was believable. Obviously, the axiomatization is again complete.

Definition 31 (Axioms) The axioms for a lying to a skeptical agent b , in case a is believed, are as follows.


$$\begin{aligned} [i_a^{\text{sc}}\varphi]B_b\psi &\leftrightarrow (B_a\neg\varphi \wedge \neg B_b\neg\varphi) \rightarrow B_b[i_a^{\text{sc}}\varphi]\psi \\ [i_a^{\text{sc}}\varphi]B_a\psi &\leftrightarrow (B_a\neg\varphi \wedge \neg B_b\neg\varphi) \rightarrow (B_a[i_a^{\text{sc}}\varphi]\psi \wedge B_a[ii_a^{\text{sc}}\varphi]\psi) \end{aligned} \quad \dashv$$

Proposition 32 The axiomatization for agent announcement to skeptical agents is complete. \dashv

The logic of agent announcements to skeptical agents can be applied to model the consecutive numbers riddle. The reader can compare the below to the first announcement in the first scenario in Section 5. We observe that both models are $\mathcal{KD}45$.



- Anne: "I know your number." *Anne is lying*



The skeptical Bill continues to be uncertain about Anne's number. He does not change his factual beliefs. But he does change his beliefs about Anne's beliefs. For example, after Anne's lie, Bill considers it possible that Anne considers it possible that he believed her. Informally, this gives Bill reason to believe that Anne has 2 and not 4. If she had 4, she would know that her lie would not be believed.

8 Lying and plausible belief

We recall the three different attitudes, presented in the introductory Section 1, towards incorporating announcements φ that contradict the beliefs $B_b \neg \varphi$ of an addressee b : (i) do it at the price of inconsistent beliefs (public announcements and agent announcements as treated in Sections 3 and 4), (ii) reject the information (announcements to skeptical agents, treated in Section 7), and (iii) accept the information by a consistency preserving process removing some old beliefs. This section is devoted to the third way. Going mad is too strong a response, not ever accepting new information seems too weak a response, we now discuss a solution in between. It involves distinguishing stronger from weaker beliefs, when revising beliefs. To achieve that, we need to give epistemic models more structure: given a set of states all considered possible by an agent, it may consider some more plausible than others, and belief in φ can then be defined as the truth of φ in the most plausible states that are considered possible. We now have more options to change beliefs. We can change the sets of states considered possible by the agent, but we can also change the relative plausibility of states within that set.

Such approaches for belief change involving plausibility have been proposed in [3, 42, 39, 9]. How to model lying with plausibility models is summarily discussed in [8, 43] (as a dialogue, these are different contributions to the same volume), and also in [9, p.54].

We continue in the same vein as in the previous sections and present epistemic actions for plausible (truthful and lying) public announcement, and for plausible (truthful, lying, and bluffing) agent announcement, by an adjustment of the epistemic actions already shown. The epistemic action for plausible public lying is the one in [8, 43, 9], the epistemic action for plausible agent lying applies the general setup of [9] to a specific (plausibility) action model. Thus it appears we can again present reduction axioms for belief change and complete axiomatizations for these logics. This is so, but it would involve not just modal belief operators $B_a \varphi$ but also conditional belief operators $B_a^\varphi \psi$ (if φ were true, agent a would believe ψ ; so that $B_a \psi$ is $B_a^\top \psi$), and additional axioms for conditional belief. We refer to the cited works by Baltag and Smets [8, 9] for further details.

Definition 33 (Plausibility epistemic models) A *plausibility epistemic model* $M = (S, \sim, <, V)$ has one more parameter than an epistemic model, namely a plausibility function $< : A \rightarrow \mathcal{P}(S \times S)$. The accessibility relations for each agent are required to be equivalence relations \sim_a . The restriction of $<_a$ to an equivalence class of \sim_a is required to

be a *prewellorder*¹⁴. ⊣

As the accessibility relations are equivalence relations, we write \sim_a instead of R_a , as before. If $s <_a t$ we say that state s is considered more plausible than state t by agent a . The transitive clause of $<_a$ is \leq_a and for $(s \leq_a t \text{ and } t \leq_a s)$ we write $s =_a t$ (*equally plausible*).


In this setting we can distinguish knowledge from belief: the agent believes φ , iff φ is true in all most plausible equivalent states; and the agent knows φ , iff φ is true in all equivalent states.¹⁵ We write $B_a\varphi$ for ‘agent a believes φ ’, as before. There is a natural way to associate an accessibility relation with such belief, based on the equivalence relations and plausibility relations, and we can then define belief in φ as the truth of φ in all accessible states, as usual. Defined in this way, these relations for belief satisfy the $\mathcal{KD45}$ properties.

Definition 34 (Accessibility relation and semantics for belief) Given an plausibility epistemic model M with \sim_a and $<_a$, the accessibility relation R_a is defined as:

$$(s, t) \in R_a \text{ iff } s \sim_a t \text{ and } t \leq_a t' \text{ for all } t' \text{ such that } s \sim_a t'$$

Given a plausibility epistemic model M and a state s in its domain, $M, s \models B_a\varphi$ iff $M, t \models \varphi$ for all t such that $(s, t) \in R_a$. ⊣

An information update that changes the knowledge of the agents, by way of changing the relations \sim_a , may also affect their beliefs, by way of changing the derived relations R_a . For example, suppose an agent a is uncertain whether p but considers it more likely that p is true than that p is false. Without reason, because in fact p is false. The epistemic state is depicted below, on the left. The arrows in the equivalence relation are dashed lines. The arrows in the accessibility relation R_a for belief are solid lines, as before. In this state B_ap is true, because p is true in the more plausible state. Now agent a is presented with hard evidence that $\neg p$ (state elimination semantics, as in truthful public announcement logic). The state where p is false is eliminated from consideration. The only remaining state has become the most plausible state. In the epistemic state depicted below on the right, $B_a\neg p$ is true. Agent a has revised her belief that p into belief that $\neg p$. In this example, belief changes but knowledge also changes. There are other examples wherein only belief changes.



Belief revision consists of changing the plausibility order between states. This induces an order between deductively closed sets of formulas. The most plausible of these is the set of formulas that are believed. This belief revision is similar to AGM belief revision, seen as

¹⁴A total, transitive, and well-founded relation

¹⁵The former represents weak belief and the latter true strong belief. There are yet other epistemic operators in this setting, safe belief, conditional belief, ... We restrict our presentation to (weak) belief.

changing the plausibility order (partial or total order, or prewellorder) between deductively closed sets of formulas. The contraction that forms part of the revision is with respect to that order. Dynamic epistemic logics for belief revision were developed to model higher order belief revision and iterated belief revision.

Just as epistemic actions, with underlying action models, generalize public announcements, plausibility epistemic actions generalize simpler forms of public plausibility updates. Accessibility relations for ‘considering an action possible’ are computed as in the case of plausibility epistemic models.

Definition 35 (Plausibility action model) A *plausibility action model* $M = (S, \approx, \prec, \text{pre})$ consists of a *domain* S of *actions*, an *accessibility function* $\approx : A \rightarrow \mathcal{P}(S \times S)$, where each \approx_a is an equivalence relation, a *plausibility function* $\prec : A \rightarrow \mathcal{P}(S \times S)$, and a *precondition function* $\text{pre} : S \rightarrow \mathcal{L}_X$, where \mathcal{L}_X is a logical language. The restriction of \prec_a to an equivalence class of \approx_a must be a prewellorder. A pointed plausibility action model is a *plausibility epistemic action*. \dashv


Definition 36 (Update with a plausibility epistemic action) The update of a plausibility epistemic state with a plausibility epistemic action is computed as the update without plausibilities (see Definition 20), except for the update of the plausibility function that is defined as:

$$(s, s) <_a (t, t) \text{ iff } s \prec_a t \text{ or } s =_a t \text{ and } s <_a t$$

where \equiv_a means equally plausible for agent a (see after Definition 33). \dashv


See [9] for details. Using these definitions, we propose the following plausibility epistemic actions for plausible (truthful and lying) public announcement that φ and for plausible (truth-telling, lying, bluffing) agent announcements that φ (by agent a to agent b). They are like the epistemic actions presented in Section 6, but enriched with plausibilities. As before, the pointed versions of these action models define the appropriate epistemic actions.

Definition 37 (Plausible (truthful and lying) public announcement) Plausible (truthful and lying) public announcements $!^{\text{pl}}\varphi$ and $!^{\text{pl}}\neg\varphi$ are the epistemic actions defined by the, respectively, right/left point of the plausibility action model depicted as follows (where $\text{pre}(!^{\text{pl}}) = \varphi$ and $\text{pre}(!^{\text{pl}}) = \neg\varphi$).



Definition 38 (Plausible agent announcement) Plausible (truthful, lying, and bluffing) agent announcements $!^{\text{pl}}_a\varphi$, $!^{\text{pl}}_a\neg\varphi$ and $!^{\text{pl}}_a\varphi$ are the epistemic actions defined by corresponding points of the plausibility action model consisting of points $!^{\text{pl}}_a$, $!^{\text{pl}}_a$, and $!^{\text{pl}}_a$, and \dashv

such that $!_a^{\text{pl}} \prec_b !_a^{\text{pl}} \prec_b !_a^{\text{pl}}$, with universal access for agent b and identity access for agent a , and with preconditions as visualized below.



Agent b 's equivalence relation is the universal relation. He cannot exclude any of the three types of announcement. Agent b 's accessibility relation for belief expresses that he considers it most plausible that a was telling the truth. It does not appear from the visualization that he considers it more plausible that a was bluffing than that she was lying. The accessibility relation is the same as in the action model for agent announcement in Section 6. Agent a 's accessibility relation is the identity on the action model. She knows whether she is lying, truthtelling, or bluffing.

As in the previous sections, an addressee rather assumes being told the truth than being told a lie or being bluffed to. But, unlike in the previous sections, we can now also encode more-than-binary preferences between actions. As lying seems worse than bluffing, we make it least plausible to interpret an announcement as lying, more plausible that it is bluffing, and most plausible that it is truthful. That is about as charitable as we can be as an addressee. This seems to be in accordance with pragmatic practice.

Proposition 39 (Axiomatization and completeness) The logics of plausible public lying and plausible agent lying have a complete axiomatization. \dashv

Proof This follows from [9, p.51]. The so-called ‘Derived Law of Action-Conditional-Belief’ is a reduction axiom for the belief postconditions of plausibility epistemic actions. This axiom involves modalities for belief, for knowledge, and conditional modalities for belief and for knowledge. The completeness proof consists of rewriting each formula in the logic to an equivalent formula in the logic of conditional belief (where conditional knowledge is definable as conditional belief).¹⁶ \square

As an example of a plausible agent announcement, we take Anne’s lying announcement “I know your number” in the first lying scenario in the consecutive numbers riddle. It suffices to refer to the depicted execution for skeptical agents in the final paragraph of Section 7. The accessibility relations for B_a and B_b are as there (after the lie, it remains the same for the speaker a , and not for addressee b). The equivalence relations for a and

¹⁶Given a plausibility epistemic model M and a state s in its domain, define R_a^φ as: $(s, t) \in R_a^\varphi$ iff $[s \sim_a t, t \leq_a t' \text{ for all } t' \text{ such that } s \sim_a t', \text{ and } M, t \models \varphi]$, and define the semantics of B_a^φ as: $M, s \models B_a^\varphi \psi$ iff $M, t \models \psi$ for all t such that $(s, t) \in R_a^\varphi$. We now have that $[_a^{\text{pl}} \varphi] B_a \varphi \leftrightarrow \neg \varphi \rightarrow B_a^\varphi [_a^{\text{pl}} \varphi] \psi$. The axioms for plausible agent lying are more complex (they also involve conditional knowledge modalities).

b are the same before and after the lie. It is a coincidence that in this example a skeptical announcement is a plausible announcement. This is a consequence of the fact that the initial belief accessibility relations are equivalence relations.

9 Conclusions and further research

Conclusions We presented various logics for an integrated treatment of lying, bluffing, and truthtelling, where these are considered epistemic actions inducing transformations of epistemic models. These logics abstract from the moral and intentional aspect of lying, and only consider fully rational agents and flawless and instantaneous information transmission. We presented versions of such logics that treat lies that contradict the beliefs of the addressee differently from those that don't, including a modelling involving plausible belief. Our main result are the various 'agent announcement' logics wherein one agent is lying to another agent and wherein both are explicitly modelled in the system.

Limitations There are limitations to our approach in view of the analysis and design of artificially intelligent agents. *Explicit agency* is missing. This lack is common in dynamic epistemic logics. We cannot distinguish agent a truthfully announcing φ from an outsider observer (or a middleman, say, if it concerns a security protocol) announcing $B_a\varphi$. Somewhat more involved, we cannot distinguish agent a lying to agent b that φ from a non-public event differently observed by a and b , but that is not seen as enacted by a . Frameworks like ATL, ATEL, and STIT have more explicit agency. *Bounded rationality* is not modelled. An agent knows if it is lying. It cannot be mistaken in the sense of announcing an inconsistency of which it is not aware. For not fully rational agents the computational cost for the liar to remember who he has been lying to, and about what, is real, and also the computational cost for the listener to check the reliability of new information. *Confidence of belief*, related to the previous, is not modelled. This requires a logic or mechanism of induction, wherein repeated observations or announcements, or by different agents, make the belief in its content stronger. In our modelling the agents do either believe the lie completely, or not at all (it abstracts from this process of increasing or decreasing confidence). *Distinguishing transmission noise from intentional noise* such as lies is not modelled, as transmission is assumed instantaneous.

Further research We see the following research directions.

Common knowledge: If agent b believes p as the result of agent a announcing p to b , we not only have B_bp but also $B_bC_{ab}p$: addressee b believes that he and the speaker a now commonly believe that p . Common belief/knowledge operators allow for more refined preconditions. A good (and possibly strongest?) precondition for agent a successfully lying that φ to agent b seems:

$$B_a\neg\varphi \wedge \neg B_b\neg\varphi \wedge C_{ab}((B_a\varphi \vee B_a\neg\varphi) \wedge \neg(B_b\varphi \vee B_b\neg\varphi))$$

In plain words, the speaker a believes $\neg\varphi$, the addressee b considers φ possible, and the announcer and addressee commonly believe/know that the announcer is knowledgeable about φ but the addressee ignorant. Common knowledge logics with dynamic modalities for information change such as action models are straightforwardly axiomatized by way of conditional common knowledge [40].

Complexity: The computational cost of lying seems a strong incentive against it. Aspects of this can also be modelled for perfectly rational agents. For example, in a different context, the computational cost of insincere voting in social choice theory [12] is intractable in well-designed voting procedures, so that sincere voting is your best strategy. The complexity of model checking and satisfiability of the logics in this paper is unclear. It seems likely that the complexity of satisfiability of truthful and lying public announcement logic is in **PSPACE**, applying similar results in [27] that builds on [21]. Also, it is unclear how to measure the complexity of a dynamic epistemic communication protocol that may involve lies.

Histories: Agents can detect lies, as in the consecutive numbers game, but we do not model that they adapt their strategies subsequently. History-based structures [31] allow the adaptation of communication strategies to the number of detected (or suspected) past lies, in the dynamic epistemic setting of perfectly rational agents.

Lying games: Consider knowledge games wherein game states are epistemic states, the agents are players making informative moves that are epistemic actions, and wherein the players' goals are knowledge of facts or other formulas true in the resulting epistemic model. Such imperfect information games (Bayesian games) with modal logic are presented in [1]. We wish to investigate *lying games* in this setting, such that each announcement can be either the truth or a lie, and where a high negative payoff is associated with the detection of a lie.

Final frontiers In multi-agent systems with several agents one may investigate how robust certain communication procedures are in the presence of few liars, or in the presence of a few lies (as in Ulam games [32]). A challenge towards philosophy is how to model a liar's paradox. In a dynamic epistemic logic this is problematic. However, promising progress towards that is reported in [26].

References

- [1] T. Ågotnes and H. van Ditmarsch. What will they say? - Public announcement games. *Synthese (Knowledge, Rationality & Action)*, 179(S.1):57–85, 2011.
- [2] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [3] G. Aucher. A combined system for update logic and belief revision. In M.W. Barley and N. Kasabov, editors, *Intelligent Agents and Multi-Agent Systems – 7th Pacific*

- Rim International Workshop on Multi-Agents (PRIMA 2004)*, pages 1–17. Springer, 2005. LNAI 3371.
- [4] G. Aucher. Consistency preservation and crazy formulas in BMS. In S. Hölldobler, C. Lutz, and H. Wansing, editors, *Logics in Artificial Intelligence, 11th European Conference, JELIA 2008. Proceedings*, pages 21–33. Springer, 2008. LNCS 5293.
 - [5] St. Augustine. De Mendacio. In P. Schaff, editor, *A Select Library of the Nicene and Post-Nicene Fathers of the Christian Church*, volume 3. Eerdmans, 1956, 1988. Translated by Rev. H. Browne. Available online at <http://www.newadvent.org/fathers/>.
 - [6] A. Baltag. A logic for suspicious players: Epistemic actions and belief updates in games. *Bulletin of Economic Research*, 54(1):1–45, 2002.
 - [7] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In I. Gilboa, editor, *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*, pages 43–56, 1998.
 - [8] A. Baltag and S. Smets. The logic of conditional doxastic actions. In K.R. Apt and R. van Rooij, editors, *New Perspectives on Games and Interaction*, Texts in Logic and Games 4, pages 9–31. Amsterdam University Press, 2008.
 - [9] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, Texts in Logic and Games 3, pages 13–60. Amsterdam University Press, 2008.
 - [10] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, Cambridge, 2001. Cambridge Tracts in Theoretical Computer Science 53.
 - [11] S. Bok. *Lying: Moral Choice in Public and Private Life*. Random House, New York, 1978.
 - [12] V. Conitzer, J. Lang, and L. Xia. How hard is it to control sequential elections via the agenda? In *IJCAI’09: Proceedings of the 21st international joint conference on Artificial intelligence*, pages 103–108. Morgan Kaufmann Publishers Inc., 2009.
 - [13] J.H. Conway, M.S. Paterson, and U.S.S.R. Moscow. A headache-causing problem. In J.K. Lenstra, editor, *Een pak met een korte broek – Papers presented to H.W. Lenstra, jr., on the occasion of the publication of his ‘Euclidische Getallenlichamen’*, Amsterdam, 1977. Private publication.
 - [14] G. d’Agostino and G. Lenzi. A note on bisimulation quantifiers and fixed points over transitive frames. *J. Log. Comput.*, 18(4):601–614, 2008.

- [15] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge MA, 1995.
- [16] H.G. Frankfurt. *On Bullshit*. Princeton University Press, 2005.
- [17] J.D. Gerbrandy and W. Groeneveld. Reasoning about information change. *Journal of Logic, Language, and Information*, 6:147–169, 1997.
- [18] U. Gneezy. Deception: The role of consequences. *American Economic Review*, 95(1):384–394, 2005.
- [19] J.L.K. Grimm and W.K. Grimm. *Kinder- und Hausmärchen*. Reimer, 1814. Volume 1 (1812) and Volume 2 (1814).
- [20] J. Hales. Refinement quantifiers for logics of belief and knowledge. Honours Thesis, University of Western Australia, 2011.
- [21] J.Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.
- [22] N. Kartik, M. Ottaviani, and F. Squintani. Credulity, lies, and costly talk. *Journal of Economic Theory*, 134:93–116, 2006.
- [23] B. Kooi. Expressivity and completeness for public update logics via reduction axioms. *Journal of Applied Non-Classical Logics*, 17(2):231–254, 2007.
- [24] B. Kooi and B. Renne. Arrow update logic. *Review of Symbolic Logic*, 4:536–559, 2011.
- [25] J.E. Littlewood. *A Mathematician’s Miscellany*. Methuen and company, 1953.
- [26] F. Liu and Y. Wang. Reasoning about agent types. Manuscript, 2012.
- [27] C. Lutz. Complexity and succinctness of public announcement logic. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 06)*, pages 137–144, 2006.
- [28] J.E. Mahon. Two definitions of lying. *Journal of Applied Philosophy*, 22(2):21–230, 2006.
- [29] J.E. Mahon. The definition of lying and deception. In E.N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, 2008. <http://plato.stanford.edu/archives/fall2008/entries/lying-definition/>.
- [30] J.-J.Ch. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science*. Cambridge Tracts in Theoretical Computer Science 41. Cambridge University Press, Cambridge, 1995.

- [31] R. Parikh and R. Ramanujam. A knowledge based semantics of messages. *Journal of Logic, Language and Information*, 12:453–467, 2003.
- [32] Andrzej Pelc. Searching games with errors—fifty years of coping with liars. *Theoretical Computer Science*, 270(1-2):71–109, 2002.
- [33] J.A. Plaza. Logics of public communications. In M.L. Emrich, M.S. Pfeifer, M. Hadzikadic, and Z.W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems: Poster Session Program*, pages 201–216. Oak Ridge National Laboratory, 1989.
- [34] H. Rott. Der wert der wahrheit. In M. Mayer, editor, *Kulturen der Lüge*, pages 7–34. Böhlau-Verlag, Köln und Weimar, 2003.
- [35] C. Sakama, M. Caminada, and A. Herzig. A logical account of lying. In *Proceedings of JELIA 2010, LNAI 6341*, pages 286–299, 2010.
- [36] F.A. Siegler. Lying. *American Philosophical Quarterly*, 3:128–136, 1966.
- [37] D. Steiner. A system for consistency preserving belief change. In *Proceedings of the ESSLLI Workshop on Rationality and Knowledge*, pages 133–144, 2006.
- [38] R. Trivers. *The Folly of Fools – the logic of deceit and self-deception in human life*. Basic Books, 2011.
- [39] J. van Benthem. Dynamic logic of belief revision. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.
- [40] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
- [41] J.H. van der Velden. *Iedereen liegt maar ik niet*. Bruna, 2011.
- [42] H. van Ditmarsch. Prolegomena to dynamic logic for belief revision. *Synthese (Knowledge, Rationality & Action)*, 147:229–275, 2005.
- [43] H. van Ditmarsch. Comments on ‘the logic of conditional doxastic actions’. In K.R. Apt and R. van Rooij, editors, *New Perspectives on Games and Interaction*, Texts in Logic and Games 4, pages 33–44. Amsterdam University Press, 2008.
- [44] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2007.
- [45] H. van Ditmarsch, J. van Eijck, F. Sietsma, and Y. Wang. On the logic of lying. In J. van Eijck and R. Verbrugge, editors, *Games, Actions and Social Software*. Springer, 2011. LNCS/FoLLI. To appear.

- [46] P. van Emde Boas, J. Groenendijk, and M. Stokhof. The Conway paradox: Its solution in an epistemic framework. In J. Groenendijk, T.M.V. Janssen, and M. Stokhof, editors, *Truth, Interpretation and Information: Selected Papers from the Third Amsterdam Colloquium*, pages 159–182. Foris Publications, Dordrecht, 1984.